

SAM (REX/Ops) Operations Meeting 2007.11.06

Present: Adam (typist), Andrew, Steve**2, Parag, Art, Krzysztof, Dehong, Randolph, Dave, Angela, Robert,

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

Still need to do retired date tests, but I've been giving Randolph other things to do. DB server v8.3.1 (with the name re-registration) is waiting on this.

From SteveW - Robert has built the linux flavor of OmniORB. Needs to do Solaris too. Robert has also built the latest Oracle Instant Client for linux (v11.?). Last we knew, it was not released for Solaris yet (in the past, Oracle makes Solaris releases).

Art: Which SAM station should people be running.

CAB & Clued0 are running 6.0.5.22srm -- this is the latest version.

CDF uses v6.0.3.9p1

MINOS uses v6.0.1.17

But the latest sam_products has v6.0.3.9p1, not the latest SAM station.

We should make a sam_products with the updated station. [**]

DØ:

Kin Yip to Adam, sam-design

Events and "complaints" :

-- fnal-farm (D0) is moving to a new station and always has a red ball in SAAG for now. We will move the farm to ignored. Can we now remove the v5 DB servers?

-- complaint about d0calib-farms on Nov. 3, 2007, not sure who started this. NGOP ? This was an automated NGOP helpdesk ticket about a network problem. Pages from the helpdesk were confused. Got e-mail that the DB calib fail over was activated, but no evidence of this in the morning. Very strange. Julie Trumbo is holding a root cause analysis meeting today.

*-- D0-9940B27.mover was OFFLINE on Nov. 2, 2007.
D71ALTO.mover and D71BLTO.mover were DEAD on Nov. 3, 2007.
The shifter Anne-Fleur has reported to enstore-admin both times.*

-- tape pro00111 became "notallowed" on 05/10/2007 .

-- samTVs:

-> samTV needed to be restarted on Nov. 2nd and 3rd.

-> clued0 seems to be empty (no picture etc.) in "sam HDTV" ?

If the web page refreshes while SAM HDTV is starting to create the new page, then the window may be empty. Steve will look into displaying a message while it is refreshing or showing the old data (with its time stamp). Steve has changed the refresh timing to minimize this effect.

-> Is there a plan to change the link of samTV on
<http://d0db-dev.fnal.gov/sam/backdoor/>
 to "sam HDTV" ?

Yes - I will do this.

Quite a few unresolved issues :

=> IT 3124: samgrid.fnal.gov scheduler not available due to disk problem ?

Machine is being replaced today.

=> IT 3122: Nov. 11, 20 MC grid jobs went into a HELD state with the error message:

=====

Held ("Globus error 7: authentication failed: GSS Major Status: Authentication Failed
 GSS Minor

Status Error Chain: init.c:497: globus_gss_assist_init_sec_context_async: Error during
 context initialization init_sec_context ")

=====

Is there any idea ? Nobody has really commented on this ticket yet.

*We'll look into it with Parag [**]. This could be due to sched being so busy that it does not reply within timeout. This will be replaced with new VDT which may help the problem.*

=>IT 3119: Joel complained about submission problem from Imperial College, UK with error like:

"Error Sending msg: Warning :Submit command diagnostic: error = 5025"

""

=> IT 3115: User at Imperial College probably met with firewall problem when sam use random ports ?? He's asking for expert help to overcome this problem.

Will look into this one -- he changed to correct DB server, but still not working. Is he running a station? SAM commands now work, but not running a project to get files. Need more information.

=> IT 3112: testCalibClient.py doesn't work (to check the healthiness of Calib DB) ?

Adam's ticket. Test script is broken because 1) dc_oracle is no longer setup by the D0 release (need to setup by hand). 2) SMT calibration detector format has changed for

Run I1b and the test program has not caught up to it. Not sure who is responsible for fixing that. And it's tied to D0 releases, which makes it hard to change.

=> IT 3110: a special sam query (for the dataset def. CSG_Skim_p20.07.01_p20.08.xx_warm_cell) seems to be very slow probably due to Oracle's inefficiency ? Any update on this ?
For this one, explain plan is using a cartesian merge join (bad if lots of rows) -- Oracle's estimate of the number of rows is very wrong.

There's also IT-3120. Querying on the file name with a wild card - explain plan is highly dependent on the details - e.g. all% is very slow. Querying on a1% is very fast. Most likely a statistics problem. Why can't Jan use other meta-data? Querying on physical stream is always fast.

Anil ran the DB re-analysis, but haven't checked it's results yet.

CDF:

Krzysztof Genser to sam-design
 CDF SAM Shift Report 10/30/07 - 11/06/07

New round of farm production in progress.
Farm is breaking number of events processed/day records.

Problems with tape IAB434 which had to have its external label replaced and a mover 9940B21 which caused unavailability of the tape IAC257.

fcdfdata099 power supply failed and another one was put in place.
 Decided to decommission the node and fcdfdata095 for spares.
These are staging nodes for dCache.

Problems with the dcap product due to scans. Looking for an optimal solution.
Lots of discussion with dCache developers. Suggested fix requires a new binary library (if called from within code) and giving a new command line option (also need a more recent client than CDF is using now). Another option is to block the scanner with IPtables (from scanning the security log, as long as not explicitly targeted at the scanner).
Robert: Are there performance/network issues from this workaround (probably leaves lots of connections).

Krzysztof is suggesting to have the default client do this secure thing. Work around is difficult to implement because of the large number of changes that would be necessary at CDF.

CDF-LTO3 library manager died several times, most noticeably on Saturday.

Thanks for the new LTO3 drives and tapes.

MINOS:

Arthur Kreymer to sam-design, minos_sam_admin
Minos status 2007 Nov 06

General -

Starting full scale Monte Carlo and reprocessing passes

Dave - can't you write to LTO3? Art - LTO3 library is not reliable yet, keep with 9940Bs for now (can migrate 9940A tapes if need more tapes - but Dave says that there are no resources for this migration).

Enstore -

Continued writing raw data archives to LTO3 (monthly .)

Net data rates about 20 MBytes/second, from a 3 GHz/GBit host.

Why is this data rate so low? This is 20 MBytes/s overall end to end rate on a good powerful machine. File is 1.7 GB. Krzysztof says that CDF typically gets 40 MB/s (which is still low).

DCache -

BlueArc -

BlueArc maintenance Nov 1 caused stale NFS handles for /grid/data mounts

There was no warning of this outside the CD Operations meeting.

No mechanism for telling anyone that this disk was going to be down.

GPFARM mounts of /minos/data and /minos/scratch are authorized !

SAM -

IT 1894 Need 'sam locate' in C++ API

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 1894: sam locate via C++ api, for simple client needs

IT 1979: Compiler warnings from sam_cpp_api

IT 2071: quadratic project startup time growth

PLAN -

Migration to 9940B tape can proceed anytime - when ?

Note: Art says that Grid certificates that originate from KCA certs are broken as of this morning. Not fixed yet.

DB:

From: Nelly Stanfield

Anil is out today, I can be contacted at x3090 during the sam mtg when & if you want to discuss the DBA part of your mtg.

Otherwise, our DBA report is below.

Oracle released the critical patch update on Oct 16.
We will be scheduling this patch on dev/int and prd databases.

We need a date for Oracle Qrtly fcdfora4~cdfopr2. 1HR
We need a date for O/S Qrtly fcdfora4~cdfopr2 2HRS

Please assign a date this qtr to apply the Oct 16 Oracle & the O/S.
We will ensure we get fcdfora1 done first depending on the date you approve.

D0 :

- COMPLEED the Oracle's critical patch on d0ofint1/d0ofprd1 on Nov 06 Tuesday

Alerts:

-Extents fo SAMPRD AP_NODE_FK_I on d0ofprd1 -- doubled

-Space wars on fcdfora4~cdfopr2 /cdf/dbs03 where the archives live.
we have been monitoring archives since july. we changed our alerts in july understanding that 83% was being used actively. but as oct we had to keep 1 less day of archives. we were average'g 500-600 per day before june but most recently we've average'g 700-800 and now as oct 30 we have been average'g slightly more then 900 archives.

we have a plan to move to a larger disks with no interruption & this action plan will be deployed later today nov 6.

the archives are 15-20M in size, and we expect to have ~400-600 due to replication on this box but the rest are due to dfc & sam activity.

on wednesday we plan to stop the stream'g to cdfofpr2 for ~30 minutes giving us chance to mine archives with no stream'g activity. we have been trying to 'mine' the archives to understand the activity but there is tons of replication activity & hard to go thru the data. we are hoping stopping replication will give us a better insight on activity on box. so more later,

Alerts:

MINOS :

Minos will be patched sometime early December. Minosdev is patched. Diana is in contact with Art regarding this schedule.

enstore/dCache:

David Berg to sam-design, enstore-admin
enstore/dcache report

Scheduled Downtimes:

- o d0en - Tue, Nov 6, 7 am to 5 pm
 - ADIC will replace Robot 2 arm; AML/2 will be down all day
 - upgrade of entire suite of enstore servers to SDE rack
 - enstore s/w updates

Downtime is underway

- o SL8500 at GCC - Thu, Nov 8
 - if we can arrange it with Sun, we will take a brief downtime on the SL8500 to switch the firmware image to a the most recent version, already loaded. This affects mostly CDF and CMS.

(was deferred from last week) Tapes are all relabeled.

Operations:

- o d0en
 - nothing to report

D0 has a lot of recycled LTO2 (1200 or so) but the write protect tabs need to be flipped. This is a manual process.

- o cdfen
 - CDF-LTO3 library manager hung and needed to be restarted, many times, with increasing frequency. *(Still investigating)*

An interesting message - received an alarm that no response from movers for previous 10 hours (when everything was working a few minutes before).

- CDF now has in production its full complement of 15 LTO3 drives
- 200 LTO3 blank tapes have been loaded into the SL8500 at GCC.

K: rate/day will be about 15/day for the next month

o cdfdca (from Alex's report)

- pool JVM restarts due to crashes; apparently it is bug in JVM code. We are changing java startup flags to avoid triggering the bug. New flags will be picked up on next java restart.
- there is slow memory growth used by pool JVMs. At some point they become stuck. It looks like this is due to high memory consumption outside java memory heap. We are changing memory allocation in startup script and it will be eventually picked up when JVM restarts.

Only seen at CDF. But perhaps not looking at the monitoring that has this information (need to look at pool restart dates). Not sure what the differences in Java are between public dCache and CDF.

o stken

- barcode replacement completed - we are ready for SL8500 firmware update.
- media changer for 9940B tapes (enstore component that communicates with the robot) became stuck and needed to be restarted. A service call has been placed for the drive that apparently triggered the problem.
- two LTO3 drives that were on loan to cdfen have not yet been returned to stken.

D0 upgrade stuff has been top priority.

o fndca

- nothing to report